IMPERFECT EQUILIBRIUM

by

Avraham Beja*

February 1990

WP # 3126-90-EFA

IMPERFECT EQUILIBRIUM

by

Avraham Beja*

February 1990

WP # 3126-90-EFA

*Visiting from The Leon Recanati Graduate School of Business Administration, Tel Aviv University

Acknowledgement: Without implicating them, I would like to thank Drew Fudenberg, Elon Kohlberg, David Kreps, Ariel Rubinstein, Dov Samet and Yair Tauman for valuable comments on earlier drafts.

Version 2/23/90

IMPERFECT EQUILIBRIUM

by

Avraham Beja

ABSTRACT

Imperfect performance equilibrium is defined as a Nash equilibrium where the players' actions may involve non-infinitesimal random errors: each player is endowed with a personal behavior pattern (his *performance function*) under which the action that he takes depends stochastically on (i) what the player sets out to do, and (ii) the potential consequences that can accrue to him from each of his possible actions. None of the feasible strategies is totally excluded, but the more costly deviations from utility maximization are less likely. An equilibrium identifies both the mutually rational *target* strategies and the *endogenously* determined error probabilities. Attractive solutions are shown to emerge in pilot studies of three well known game models - the coordination game, the repeated prisoner's dilemma, and the chain store problem. In particular, equilibrium behavior exhibits an appealing sensitivity to *how desirable* the key outcome is, e.g. cooperation is more likely to arise when the fruits of cooperation are higher. Unlike earlier studies dealing with similar issues, the results for imperfect equilibrium do not depend on asymmetric information or on lengthy repetitions of the game under consideration.

VERSION 2/23/90

IMPERFECT EQUILIBRIUM

by

Avraham Beja

I. INTRODUCTION.

The continuing effort to improve and refine the prevalent solution concepts for game models of economic behavior attests to the many important issues that remain unresolved. This note offers a new, or rather 'modified', solution concept for strategic games which is based on a number of familiar ideas. The new arrangement of the familiar building blocks can help resolve a variety of puzzling problems, and we use three well known game models - the coordination game, the repeated prisoner's dilemma, and the chain-store problem - as pilot applications. The examples demonstrate the power of the proposed approach to identify solutions that are in some respects more attractive than those offered by the current theories. In particular, the results do not depend on asymmetric information or on lengthy repetitions of the games under consideration.

The proposed solution concept, termed "imperfect performance equilibrium", is based on the idea that players may fail to do what is apparently best for them. In this respect, our approach is clearly related to a number of earlier models, and particularly to Selten's (1975) notion of "trembling hand" perfect equilibrium (hence the association in the terminology). "Perfect" equilibria take this potential failure into account only as a remote and extremely unlikely possibility, which can affect one's decisions only in the choice between actions which would otherwise seem equally optimal. In the approach proposed here, on the other hand, uncertainty about the agents' ability to actually implement their optimal strategies is considered an essential aspect of the underlying environment - an aspect that *always* affects the agents' behavior and must therefore be made an *explicit* part of the game model. For one thing, what the player actually *does* need not always coincide with what he sets out to do. A typical situation is when the "player" in the model represents a whole organization composed of many individuals whose diverse interests do not always coincide with the stated interests of the organization. Without explicitly formalizing the implementation process, it is clearly conceivable that the agents entrusted with the execution of a prescribed strategy

may in fact implement an alternative that better suits their own self interests. Secondly, failure to take the best strategy may also be due to inaccurate "computation", e.g. the players' assessments of the consequences of alternative strategies may involve a random approximation error.

In the proposed approach, the game model is extended to include a specification of the players' abilities to implement their desired strategies. Each player is endowed with a personal behavior pattern, termed his "performance function", under which the action that he takes depends stochastically on (i) the target strategy that he would rationally like to adopt, and (ii) the potential consequences that can accrue to him from each of his feasible strategies. Our analysis assumes that the players will be careful to avoid the costlier deviations from optimal behavior more than the relatively trivial ones (this is very much in the spirit of Myerson's "proper equilibrium", 1978). In the organizational setting, agents are expected to be more reluctant to implement an unauthorized alternative that greatly harms the organization, since in the long run this will also harm them. From the computational viewpoint, smaller errors are sufficient to induce adoption of an alternative with a low opportunity loss, while only major errors can induce adoption of an alternative with a high opportunity loss, hence deviations of the latter type are again less likely.

An "imperfect equilibrium" identifies mutually consistent strategies of sophisticated players who know that deviations from value maximizing decisions can indeed occur: they actively adjust their decisions to take advantage of the potential errors of others (while at the same time taking into account that the decisions they adopt may be inaccurately implemented). This is in essence a Nash equilibrium in a limited and especially structured domain of "imperfectly executed" strategies. The effect can sometimes be a further "refinement" of the results offered by the familiar Nash equilibrium concepts. But imperfect equilibrium can also bring forth quite different solutions, in which the superficially "optimal" actions are superceded by attractive strategies that would have seemed "irrational" if errors were inconceivable.

Imperfect equilibrium exhibits yet another interesting feature - it can account for different behavior in games whose basic structures seem indistinguishable under the prevalent theories. In other words, a class of apparently similar games possessing essentially the same equilibria (Nash, perfect, or proper) can exhibit in the imperfect equilibrium framework widely diverging solutions, depending on whether the payoffs in some contingencies are "very high" or just "high". This occurs because (1) the extent to which a player may be inclined to adjust his choice of strategy due to the presumed potential errors of others may be quite sensitive to the *magnitudes* of the differences between the payoffs associated with alternative strategies (not just the *signs* of these differences), and (2) the propensity of the opposition to make errors may itself be affected by the magnitudes of the opportunity losses involved.

This short note is suggestive rather than definitive. After formalizing the idea of "imperfect performance", imperfect-performance-equilibrium is defined and shown to exist for all n-person strategic games with finitely many strategies. Some examples are then offered to demonstrate different applications. In the first example, imperfect equilibrium acts as a "refinement" by singling out the more attractive and more plausible outcome out of two Nash equilibria that are both (trembling hand) "perfect". The second example deals with a class of prisoner-dilemma games, and demonstrates that when the fruits of cooperation are sufficiently high it is logical to expect a high probability of cooperation in the first stage even if the game is only played twice. In the last example, the limited performance capabilities of two potential entrants to a monopoly market gives rise to seemingly self-defeating (but in fact rational) tough action on the part of the incumbent monopolist.The presentation ends with a few concluding remarks.

II. IMPERFECT PERFORMANCE

We start with a formal description of the behavior of agents in an environment where perfect performance is only an ideal reference that cannot be expected to take effect in practice. Suppose that a player (named 'j') has K pure strategies available to him. The probabilities for

any of these strategies to be in fact implemented by the player will be assumed to depend on three factors: (i) what the player sets out to do, (ii) the potential consequences associated with each of his feasible strategies, and (iii) the player's "capability". Suppose that $p^j \in \Re^K$ denotes the player's "target" strategy (i.e. p^j_k is his target probability for using his k^{th} pure strategy) and that $x^j \in \Re^K$ denotes the consequences that accrue to him from his various strategies (i.e. x^j_k is his utility if he in fact implements his k^{th} pure strategy). The effective mixed strategy that describes the actual behavior of player j is then given by $f^j(p^j,x^j) \in \Re^K$ (i.e. $f^j(p^j,x^j)_k$ is the probability that player j will in fact use his k^{th} pure strategy). The function f^j reflects the (limited) capability of player j, and will be termed his **performance function**. The performance function is to be considered a basic attribute of each agent, describing his behavior in a variety of potential situations (possibly involving different numbers of pure strategies). Our assumptions with respect to the player's behavior will be stated as assumptions about the performance function f. This is formalized as follows.

Let X denote the set of all finite tuples of real numbers (i.e. $X=\bigcup_{K=1}^{\infty}\Re^K$). For every $x \in X$ let $K(x)$ denote the number of elements in x (reflecting the number of pure strategies available to the player in a game under consideration). Also, let P denote the corresponding set of all mixed strategies, i.e. $P=\{p \in X: p_i \geq 0, \Sigma p_i=1\}$. The performance function is a mapping f with domain $D=\{(p,x) \in P \times X : K(p)=K(x)\}$ and range in P, satisfying $K(f(p,x))=K(p)$. Although this formulation is designed to distinguish between the target strategy and the implemented strategy, it does admit as a special case the *perfect implementation* assumed in the classical analysis, which is represented in our model by a function f^* satisfying $f^*(p,x)=p$. One can of course also think of performance patterns that are "almost perfect", and we shall occasionally use $\rho(f)=\sup_{p,x,k} |f(p,x)_k-p_k|$ as a measure of the "imperfection" associated with a performance function f.

The following six properties represent "reasonable" attributes that an imperfect performance function may or may not satisfy (our subsequent results depend on various combinations of these

properties). We start with a straightforward continuity assumption, which states that in situations which are barely distinguishable the agent's behavior is essentially "similar".

(1) $f(p,x)$ is continuous in p and x.

Unlike the earlier "trembling hand" models, this study does not restrict the deviations from the target strategy to be necessarily "infinitesimal". To maintain a reasonably tight association between the target strategies and the implemented strategies, we invoke instead two weaker assumptions. The first of these states that the best way to avoid using a given pure strategy is by simply not attempting to use it (rather than by "sophisticated" manipulation of the target probabilities for this or for other pure strategies).

(2) The target strategy p minimizes $f(p,x)_k$ if and only if $p_k=0$.

Whereas assumption (2) posits a (limited) monotonicity of the actual probability with the target probability for a fixed k, assumption (3) deals with comparisons across different k for a fixed p.

(3) If $p_k>p_i$ and $x_k \geq x_i$ then $f(p,x)_k>f(p,x)_i$.

The next assumption reflects the familiar idea of the"trembling hand": when a player's performance is short of perfect, any of his pure strategies has some chance, perhaps small, to be the one that is actually used.

(4) For $f \neq f^*$, $f(p,x)_k>0$ for all p,x, and k.

The last two assumptions reflect the other familiar idea that a strategy is less likely to be implemented if its payoff is (relatively) lower. Assumption (6) below relates only to the unintentional use of pure strategies involving an opportunity loss. To identify these, let $\mu(x)$ denote the value of the largest element in x, i.e. $\mu(x)=\max\{x_i: i=1,...,K(x)\}$. We assume

(5) $f(p,x)_k$ is non-decreasing in x_k and non-increasing in x_i for $i \neq k$.

(6) If $p_k=p_i=0$ and $0<\mu(x)-x_k<\mu(y)-y_i$, then $f(p,x)_k>f(p,y)_i$.

<u>Remarks</u>:

1. Assumptions (5)-(6) (and to some extent also (3)) attribute some "rationality" to the deviations of the implemented strategies from the corresponding target strategies. In our model, this

(limited) rationality in implementation does not replace, but rather complements, the rationality of the target strategies (the choice of the target strategies, rational or otherwise, is an issue distinct from the implementation of a strategy, once it is considered to be the "target"; the standard rationality assumptions regarding the target strategies will be stated in due course in the definition of equilibrium).

2. Assumptions (1)-(2)-(3) and (5) are direct generalizations of the classical perfect performance, in that they are clearly satisfied by the perfect performance function f*. Technically speaking, this applies (trivially) also to assumption (4), as stated. On the other hand, f* clearly does not satisfy the given statement of assumption (6). This difference between the formulations of (4) and (6) is designed to emphasize that assumption (6), which is used formally only in section V, applies only in the context of strictly imperfect performance.

3. Assumption (6) is the exception relative to the other five assumptions also in that it relates to "opportunity losses", defined as *differences* in utility payoffs.

4. Assumption (2) is the mildest monotonicity property under which the target "best responses" are always the same as those of an agent with perfect performance.

5. The utility numbers assigned to the various outcomes are of course non-unique representations of the player's preferences. Since the performance function specifies the effective strategies associated with alternative utility realizations in X, this necessarily presumes that the utility scale is set before the function f is defined. Also, for our assumptions to apply across different games involving different outcomes, one common utility scale must be fixed for all the games involved before any cross comparisons are made.

III. EQUILIBRIUM IN STRATEGIC GAMES.

To allow for imperfect performance in strategic games, the formal description of the game must be extended to include a statement of the players' capabilities. A *performance-dependent n-person strategic game* will be defined by

$$\Gamma=(S^1,\ldots,S^n;\ U^1,\ldots,U^n;\ f^1,\ldots,f^n)$$

where each S^i is a non-empty finite set of "pure strategies", U^i is a real valued "utility-payoff" function defined on $S^1 \times S^2 \times ... \times S^n$, and f^i is the i^{th} player's performance function. We let the elements of each S^i be consecutively numbered and renamed as $\{1,2,...,|S^i|\}$. A mixed strategy, or in short **strategy**, for player i is then an element $p^i \in P$ such that $K(p^i)=|S^i|$ (this includes the degenerate mixed strategies where one of the elements of p^i is the unit mass and all other elements have zero values). An n-tuple $\underline{p}=(p^1,...,p^n)$ of (mixed) strategies for each of the players is termed a **strategy combination**, and the set of all strategy combinations is denoted by M. The utility function is extended to mixed strategy combinations in the usual way, i.e.

$$U^i(\underline{p})=U^i(p^1,...,p^n)=\Sigma_{k_1=1}^{|S^1|}\Sigma_{k_2=1}^{|S^2|}...\Sigma_{k_n=1}^{|S^n|}(p^1)_{k_1}(p^2)_{k_2}...(p^n)_{k_n}U(k_1,...,k_n)$$

In the imperfect-performance context, the players' expected utilities are obtained by applying the extended utility function to the effective strategy combination (incorporating the actual implementation probabilities), not to the target strategy combination. When a player makes his strategy choice, his decision and his behavior depend on the potential, or tentative, payoffs that can accrue to him, given the other players' effective strategies, if he in fact uses any one of his pure strategies. This will be represented as follows. For $i=1,...,n$ and $k=1,...,|S^i|$ let $e^i_k \in P$ satisfy $K(e^i_k)=|S^i|$, with $(e^i_k)_k=1$ and $(e^i_k)_j=0$ for $j \neq k$, and let $V^i(e^i_k|\,\underline{p})$ be defined by

$$V^i(e^i_k|\,p^1,...,p^n)=U^i(p^1,...,p^{i-1},e^i_k,p^{i+1},...,p^n).$$

We can now define a modified notion of Nash equilibrium for strategic games with imperfect performance.

<u>Definition</u>. An **imperfect performance equilibrium**, or in short **imperfect equilibrium**, for a game Γ is a tuple of elements $(\langle p^1,...,p^n\rangle;\langle q^1,...,q^n\rangle;\langle x^1,...,x^n\rangle)$, where $\langle p^1,...,p^n\rangle$ and $\langle q^1,...,q^n\rangle$ are strategy combinations, and where, for $i=1,...,n$, $x^i \in X$ with $K(x^i)=|S^i|$, such that

(1) $p^i \in \arg\max_\xi \Sigma_k f^i(\xi,x^i)_k (x^i)_k$

(2) $q^i=f^i(p^i,x^i)$

(3) $(x^i)_k=V^i(e^i_k|\,q^1,...,q^n)$.

The strategy combination $<p^1,...,p^n>$ will be termed the **target (strategy) equilibrium** associated with the above imperfect performance equilibrium, and $<q^1,...,q^n>$ will be the corresponding **effective (strategy) equilibrium**.

Theorem. If the performance functions $f^1,...,f^n$ satisfy assumptions (1) and (2) then the performance-dependent game $\Gamma=(S^1,...,S^n; U^1,...,U^n; f^1,...,f^n)$ has at least one imperfect performance equilibrium.

Proof (outline).

Let $\chi=\{x^*=(x^1,...,x^n): x^i\in X, K(x^i)=|S^i|$, and for $k=1,..,|S^i|$ $\min_{\underline{p}}V^i(e^i{}_k|\,\underline{p})\leq(x^i)_k\leq\max_{\underline{p}}V^i(e^i{}_k|\,\underline{p})\}$.

For every $(\underline{p},\underline{q},x^*)\in M\times M\times\chi$ let $T(\underline{p},\underline{q},x^*)$ be the set of tuples $(\underline{s},\underline{r},y^*)\in M\times M\times\chi$ satisfying for $i=1,...,n$

1. $s^i\in \arg\max_\xi\Sigma_k f^i(\xi,x^i)_k(x^i)_k$
2. $r^i=f^i(p^i,x^i)$
3. For $k=1,...,|S_i|$, $y^i{}_k=V^i(e^i{}_k|\,r^1,...,r^n)$

The proof consists of the following points.

A. The domain of T is a compact convex subset of finite dimensional Euclidean space.

B. Assumption (2) implies that $s^i\in \arg\max_\xi\Sigma_k f^i(\xi,x^i)_k(x^i)_k$ if and only if

$s^i\in \arg\max_\xi\Sigma_k (\xi)_k(x^i)_k$, hence $T(\underline{p},\underline{q},x^*)$ is closed and convex.

C. Assumption (1) then implies that T is upper semicontinuous.

Thus T satisfies the conditions of the Kakutani fixed point theorem, and every fixed point of T defines an imperfect performance equilibrium.

Remark. A strategy combination $\underline{p}$ is a *trembling-hand-perfect* equilibrium of the conventional n-person game $\Gamma^o=(S^1,...,S^n; U^1,...,U^n)$ [which is of course equivalent to the performance-dependent game $\Gamma^*=(S^1,...,S^n; U^1,...,U^n; f^*,...,f^*)$] if and only if $\underline{p}$ is a limit of a sequence $\{\underline{p}_t, t=1,2,...\}$ of target equilibria for the games $\Gamma_t=(S^1,...,S^n; U^1,...,U^n; f^1{}_t,...,f^n{}_t)$ where, for all i and t, $f^i{}_t\neq f^*$ also satisfy (4) and $\rho(f^i{}_t)\rightarrow 0$ (this observation follows directly from

Myerson's (1978) definition of ε-perfect equilibria). Similarly, the *proper* equilibria of the conventional game constitute the subset of these limits where all f^i_t belong to a special class F°, where each f∈ F° satisfies the condition

$$\text{If } x_i<x_k \text{ (and } p_i\leq p_k) \text{ then } f(p,x)_i<\rho(f)f(p,x)_k.$$

This demonstrates the close relationship between imperfect equilibria and perfect or proper equilibria.

IV. APPLICATION: IMPLICIT COORDINATION.

Consider the familiar class of implicit coordination games represented by the following matrix, where H≥0.

	a	b
a	(H,H)	(0,0)
b	(0,0)	(1,1)

The two players are (equally) awarded if their actions are "coordinated", but they cannot communicate with each other and must act independently. Both (a,a) and (b,b) are Nash equilibria, and for all H>0 they both also qualify as "perfect" and "proper". This is quite intuitive whenever H is not too different from 1: with H>1, both players would prefer (a,a), but if one player has *some* reason to believe that the other's action will be b then he must himself take action b, and since the argument is symmetrical this only reinforces his belief that the other will indeed choose b, and so forth. But this argument seems much less intuitive when H is "very large": even if for some reason (possibly associated with past history) the action b is a "focal point", when H is in the order of magnitude of, say, 10^{12} then each one of the two players should be confident enough to argue that the overwhelming advantage of (a,a) must be sufficient to drive the other from the focal point, and one can reasonably expect that the players will indeed act accordingly. As H gets smaller and smaller, (a,a) becomes less and less attractive and (b,b) becomes more and

more attractive. When the value of H is almost indistinguishable from zero the argument against (a,a) is essentially equivalent, indeed almost identical, to the arguments that have been offered extensively in the context of H=0 to demonstrate the unacceptability of Nash equilibria that are not also "perfect". It follows that a desirable solution concept should rule out the poor outcome whenever H is "very large" or "very small", but admit both (a,a) and (b,b) not only when H=1 but also whenever H is neither very large nor very small.

<u>Proposition</u>. Let $\Gamma(H)$ denote the above implicit coordination game with a specific value of H and performance functions f^1 and f^2, and suppose that f^1 and f^2 satisfy assumptions (1)-(5) and that $\rho(f^1)+\rho(f^2)=\varepsilon>0$, however small. If H is sufficiently large then the *unique* target equilibrium for $\Gamma(H)$ is $((1,0),(1,0))$, if H is sufficiently small the *unique* target equilibrium for $\Gamma(H)$ is $((0,1),(0,1))$, and if H is sufficiently close to one then both $((1,0),(1,0))$ and $((0,1),(0,1))$ are target equilibria for $\Gamma(H)$.

The proposition is very intuitive, but the fact that the equilibrium error probabilities emerge endogenously requires that a few details be attended to. A detailed proof is given in appendix A.

Compare the stated proposition to the view expressed by Aumann and Sorin (1989), who assert that "without repetition there is no hope of getting the kind of result (that ensures cooperation)...people may perhaps *learn* to cooperate, but they cannot be expected to do so in a one-time, static situation" (p.8). In contrast, our proposition suggests that in games with common interests (defined by Aumann and Sorin as games in which there is one payoff pair that strongly Pareto dominates all other payoff pairs) cooperation emerges naturally if the mutually preferred outcome is sufficiently desirable (i.e. "worth taking chances for"). To aim at the desired outcome, a player need not really believe that the other player will necessarily also do the same, it suffices that he believe that the other *might* do so, even if by "mistake".

Furthermore, the stated proposition admits a straightforward generalization to more than two players. As would be expected on the basis of simple common sense, the minimal

value of the target payoff required to guarantee cooperation by many players increases with the number of parties involved in the process.

V. APPLICATION: REITERATED PRISONER'S DILEMMA GAME

Our second application of the imperfect equilibrium concept deals with the ever-intriguing "prisoner's dilemma" games, represented here by the following matrix, where 'c' stands for "cooperate", 'd' stands for "defect", and H>1.

	c	d
c	(H,H)	(0,H+1)
d	(H+1,0)	(1,1)

For both players, defection strongly dominates cooperation, and this leads inevitably to the Pareto inferior outcome (1,1). The greater paradoxes emerge when this game is repeated more than once. Persistent defection does *not* dominate all other strategies, but the conclusions of the classical analysis are no less disturbing. If the number of repetitions is predetermined, then backwards induction, equivalently the sequential elimination of dominated strategies, again singles out persistent defection by both players as the unique equilibrium. Even if there is always a very high probability that repetitions will continue and there is almost no discounting, persistent defection by both players is still considered by the theory a reasonable solution for two perfectly rational agents (besides the other, more attractive, solutions that are now also admissible). Perhaps most disturbing, and certainly most relevant for the present analysis, is the observation that the above conclusions are totally independent of the magnitude of the payoff H, whereas both intuition and casual observation of economic and political systems suggest that as the fruits of cooperation grow larger so does the tendency to cooperate.

For a single iteration of the prisoner's dilemma game, imperfect equilibrium maintains the essence of the classical results. Regardless of any potential errors in strategy

implementation on either side, agents still attempt to minimize their use of the dominated strategy 'c' and hence their target probability of cooperation is zero. But for the repeated game, imperfect equilibrium exhibits an appealing departure from the less attractive classical results even when there are only *two* iterations.

Proposition. Let $\Gamma(H)$ denote the repeated game consisting of two iterations of the above prisoner's dilemma game with a specific value of H and performance functions f^1 and f^2, and suppose that f^1 and f^2 satisfy assumptions (1)-(4) and (6). As H gets large, the effective probability of cooperation in the first iteration for each of the two players in an imperfect performance equilibrium for $\Gamma(H)$ approaches 1/2.

A detailed formal proof of the proposition is given in appendix B. It may be instructive, however, to also review informally how this apparently surprising result emerges. Consider player 'one'. It is clearly in his best interests to defect in the second stage, no matter what player 'two' does in the first stage: if he cooperates rather than defects, he incurs an opportunity loss of one payoff unit. Suppose, however, that 'one' adopts by mistake a strategy that imitates at stage two what 'two' does at stage one. This will cause 'one' an opportunity loss of one unit if and only if 'two' cooperates at stage one. The expected loss to 'one' due to this mistake is thus exactly the probability that 'two' cooperates at stage one, say $Q^2{}_c$. Similarly, if 'one' uses by mistake a strategy that responds in the opposite fashion by playing at stage two what 'two' did *not* play at stage one, he incurs an expected opportunity loss of $Q^2{}_d=1-Q^2{}_c$. If $Q^2{}_c$ is low, the expected loss to 'one' due to the latter type of mistake is higher, and therefore his mistakes will be biased towards the "tit-for-tat"-like responses. When H is very high, this induces player 'two' to cooperate in the first stage, in the (remote) hope of being highly rewarded in the second stage if 'one' happens to make the (more likely) mistake of responding in kind. Hence a low $Q^2{}_c$ cannot apply in equilibrium. But certainly a high $Q^2{}_c$ cannot apply in equilibrium, because if player 'one' mistakenly lets his action at stage two depend on what 'two' has done, then his opportunity losses will make it more likely for him to respond with cooperation if 'two' defects,

in which case there is certainly no incentive for 'two' to cooperate at stage one. This argument of course applies for both players. It follows that, in equilibrium, cooperation and defection must be almost equally likely at stage one if H is sufficiently high.

The idea that "bounded rationality" can induce a rational motive for cooperation in the repeated prisoner's dilemma game is of course not new. The line of research that seems most closely related to our proposition is probably that originated by Kreps, Milgrom, Roberts and Wilson (1982), who showed that cooperation in the first stages (in a sequence of many repetitions) emerges rationally if there is some probability that one of the players is for some reason committed to playing tit-for-tat. The key to their result is the arbitrary selection of tit-for-tat as the only deviation from "rationality" that is anticipated by the opposition. Put differently, the result depends on an exogenous assumption that tit-for-tat is *more likely than other conceivable deviations* from utility-maximizing behavior.[1] In imperfect-performance-equilibrium, the balance between the relative probabilities of alternative deviations from uniform defection is determined *endogenously* by the structure of the associated payoffs, and this balance gives rise to cooperation.[2] Like in Kreps-Milgrom-Roberts-Wilson and in Aumann-Sorin, the proposition on imperfect equilibrium identifies conditions under which cooperation is *guaranteed* to emerge in equilibrium, rather than being simply possible.[3] Furthermore, in all the familiar studies that consider bounded rationality, cooperation emerges only if very many future repetitions are

[1] A similar difficulty applies to many variants and extensions of this approach, see e.g. Fudenberg and Maskin (1986) and the references cited there.

[2] This complies with the call issued by Aumann and Sorin (1989) in their discussion of the Kreps-Milgrom-Roberts-Wilson approach, where they say that "...one would have liked a stronger result, in which one is led to a cooperative outcome entirely endogenously. For example, this would be the case if the perturbation were a mixture of *all* alternative strategies..." (p.11).

[3] *Possible* cooperation (but not *guaranteed* cooperation) can easily emerge in equilibrium for infinite repetitions of the prisoner's dilemma. Possible cooperation can also emerge in fintely many repetitions of the game, e.g. if the players have only bounded capabilities for using complex strategies (as in Neyman, 1985).

anticipated, whereas the result for imperfect-performance-equilibrium applies even if the number of repetitions is as small as two.

VI. THE CHAIN-STORE PARADOX

For one further indication of the variety of potential applications, we give in this section a brief *informal* discussion of a model involving an *asymmetric* conflict of interests, with essentially no cooperation. The model is a modified "short" version of Selten's (1977) celebrated chain-store paradox. Suppose that there is one 'incumbent' operating in two markets, and two potential 'entrants', one for each of the two markets, who act sequentially in time. If a potential entrant stays out ('o'), he gets 0 and the incumbent gets H; if he enters ('e') and the incumbent succumbs ('s') the entrant gets 1 and the incumbent gets 1, and if he enters and the incumbent plays tough ('t') he suffers a loss and gets -L while the incumbent gets 0. The payoffs to the incumbent in the two markets are additive. It is clear that with perfect performance the only subgame-perfect equilibrium for this game has the entrants enter and the incumbent succumb in both markets (this equilibrium is of course also "trembling-hand-perfect" and "proper"). In particular, this holds no matter how high the payoff H to the incumbent if the entrant stays out, or the loss L to an entrant if he enters and the incumbent fights. The greater paradox here is that a similar result applies even if the number of markets is very large: with perfect performance, there is no credible way for the incumbent to try to deter entry by playing tough. Could there be any *indirect* reason for the incumbent to play tough and, in particular, could such a reason apply even if the number of markets is small?

One appealing approach for resolving the chain-store paradox has been to assume some uncertainty about the behavior of the *incumbent* who faces a large number of potential entrants. Such an incumbent can capitalize on the fears of the potential entrants that he may be a "type" that does not maximize his own utility but always plays tough, and deliberately reinforce these

fears by pretending to be one.[1] To emphasize the difference between this approach and the forces that determine the imperfect-performance-equilibrium solution, we not only let the number of markets be the smallest one that can involve entry-deterring responses, but also consider the extreme case where the performance capability of the incumbent is perfect (and known to the entrants to be such). The analysis leads to the following conclusion: when the second entrant's performance capability is less then perfect, and H is sufficiently high, it pays for the incumbent to play tough against the first entrant with probability which is almost as high as the probability of succumbing to him. When L is also high, this can deter both potential entrants from being the first to enter, and the incumbent prevails.

To see how this happens, consider the decision of the incumbent when the first entrant enters. Everyone knows that if and when the second entrant enters the incumbent must succumb (being a perfect utility-maximizer). Hence it pays the second potential entrant to enter, no matter how the incumbent reacts to the first entry (call this strategy 'e,e'). But the second entrant's performance is less than perfect, and he may make a mistake and adopt another, non-optimal, strategy. What are the strategies available to the second entrant, given that the first entrant has entered? Besides the optimal strategy 'e,e', the second entrant has three other strategies: staying out, no matter what ('o,o'), entering if the incumbent succumbs to the first entrant and staying out if the incumbent plays tough with the first entrant ('e,o'), and the opposite strategy of staying out if the incumbent succumbs and entering if the incumbent plays tough ('o,e'). Given that the incumbent will play optimally at the last stage, the (expected) opportunity loss to the second entrant from using 'e,o' is exactly the probability that the incumbent will play tough to the first entrant, say $Q^1{}_t$, and the opportunity loss associated with 'o,e' is the complement probability, say $Q^1{}_s=1-Q^1{}_t$. If the incumbent were to succumb to the first entrant with very high probability, then the opportunity loss to the second entrant from using 'o,e' would be higher

---

[1] See for example Kreps and Wilson (1982), Milgrom and Roberts (1982) and Fudenberg and Levine (1989).

than the opportunity loss from using 'e,o', and hence 'e,o' would be a more likely strategy for him to adopt (by mistake) than 'o,e'. But if this is the case, and H is sufficiently high, it pays for the incumbent to play tough to the first player, incurring a small immediate loss and a small expected future loss (in case the second entrant is using 'o,e') for the sake of the larger expected gain in the more likely event that the second entrant is in fact using 'e,o'. For large H, equilibrium obtains *only if* $Q^1{}_t$ is almost equal to $Q^1{}_s$. Since this is the optimal response of the incumbent to the first entry (given the anticipated behavior of the second entrant) it does not pay to enter first if L is (substantially) more than 1. In this analysis, the perfect performance of the incumbent simplifies the exposition, but it is not essential for the conclusion.

A remarkable feature common to the above example and the two previous ones is that the players' attributes were assumed to be common knowledge. In the framework of these models there was thus no room for "pretending" and no need for "learning" (indeed, there was also not too much time for either). The possibility of errors, some of which being for good reason more likely than others, can account for much of what is often being ascribed to more complex aspects, such as the sequential evolution of asymmetric information.

VII. CONCLUDING COMMENTS

Imperfect-performance-equilibrium is a solution concept for the extended model that we have labeled "performance-dependent-game". The game model itself highlights the basic premise of our suggested solution: for many interesting problems, it is not sufficient to consider only the players' available strategies and their preferences with respect to potential outcomes - it is essential that their "performance capabilities" be also considered explicitly. A similar view is of course implicit also in other approaches, e.g. the studies (like Neyman, 1985, Gilboa and Samet, 1989, etc.) that take explicit account of the players' abilities to handle complex strategies.[1]

[1] Another extension of the game model is used by Rubinstein (1986), where the players' strategy sets are augmented to include a choice betwen alternative automatons with different "capabilities".

Imperfect performance equilibrium is of course not insensitive to the inclusion or exclusion of dominated strategies, and the solutions that it offers might therefore seem to depend on "irrelevant alternatives". But under the underlying premise of imperfect equilibrium, feasible alternatives are not necessarily irrelevant even if they are strongly dominated. This may raise some questions of general modeling philosophy, because it is clearly impossible to include explicitly in models of real-world phenomena *all* the actions which are technically feasible for the players, regardless of how irrational these actions may seem. But the potential for modelling interesting economic and social problems as performance-dependent-games is not destroyed, because "extremely inferior" alternatives are so unlikely to surface that they can be considered essentially "irrelevant".

Selten's (1975) discussion of the "trembling hand" in the context of extensive-form games does not say much about the origin or nature of the trembles. Although the error probabilities at the various information sets in each of his "perturbed games" are assumed mutually independent, the *sequences* of perturbed games (whose equilibria converge to perfect equilibria of the initial game) may involve implicit subtle relationships between the error probabilities at the various stages. Imperfect equilibrium takes the explicit view that errors are expressed by deviations from a target strategy to another *strategy*, and the effects of these errors in any given performance-dependent-game therefore exhibit a degree of serial dependence.

In the examples offered in this exploratory study, a sufficiently high value of the payoff H was sufficient to single out outcomes that contrast the prevalent theories, and this applied even in games which are very short. It is almost irresistably tempting to conjecture what may happen to the imperfect performance equilibria if these games are extended to much longer repetitions. For the repeated prisoner's dilemma game, it is natural to expect a robust Kreps-Milgrom-Roberts-Wilson type of result, where the "tit-for-tat"-like strategies emerge endogenously as the more probable deviations from strict utility maximization. For the repeated coordination game, one can expect an outcome similar to the result of Aumann and Sorin (1989)

on "cooperation and bounded recall". In Aumann and Sorin's analysis, the possible deviations from rational behavior are exogenously assumed to consist of all "bounded-recall strategies". It would be extremely attractive if the desired bounded recall strategies emerge endogenously in the imperfect equilibrium context (by merit of the associated payoffs) as much more likely than alternative strategies that never forget-and-forgive. In all of the examples we have considered,, the minimum value of H required to single out the desired outcome can be expected to decrease as the number of repetitions gets large. Of course, one cannot hope to analyze these complex models successfully with the simple methods used here.

Finally, our discussion of imperfect performance equilibrium in the restricted context of strategic games in normal form probably suppresses interesting issues associated with "imperfect sequential equilibria" in extensive-form game models. Here, too, the scope for further research seems most promising.

APPENDIX A: PROOF OF PROPOSITION FOR THE COORDINATION GAME.

Without loss of generality, suppose $f^1 \neq f^*$.

(a) Define $\alpha = f^1((0,1),(0,1))_1$.

Consider $H > \alpha^{-1}$, and suppose $(\underline{p},\underline{q},x^*)$ is an imperfect equilibrium for $\Gamma(H)$.

Then $x^1{}_1 = Hq^2{}_1 \geq 0$ and $x^1{}_2 = q^2{}_2 \leq 1$, hence (5) implies $f^1((0,1),x^1)_1 \geq \alpha$, and then (2) implies $f(p^1,x^1)_1 = q^1{}_1 \geq \alpha$, hence $x^2{}_1 = Hq^1{}_1 \geq H\alpha > 1 > q^1{}_2 = x^2{}_2$. Hence by (2) $p^2 = (1,0)$.

Then (3) implies $q^2{}_1 > q^2{}_2$, hence $x^1{}_1 = Hq^2{}_1 > q^2{}_1 > q^2{}_2 = x^1{}_2$.

Hence by (2) $p^1 = (1,0)$. I.e., the unique target equilibrium for $H > \alpha^{-1}$ is $((1,0),(1,0))$.

(b) Similarly define $\beta = f^1((1,0),(1,0))_2$ and consider $0 \leq H < \beta$.

If $(\underline{p},\underline{q},x^*)$ is an imperfect equilibrium for $\Gamma(H)$ then $x^1{}_1 = Hq^2{}_1 < 1$ and $x^1{}_2 = q^2{}_2 \geq 0$, hence by (2) and(5) $q^1{}_2 = f^1(p^1,x^1)_2 \geq f^1((1,0),x^1)_2 \geq \beta$, hence $x^2{}_1 = Hq^1{}_1 < H < \beta \leq q^1{}_2 = x^2{}_2$, hence by (2) $p^2 = (0,1)$.

Now (3) implies $q^2{}_1 < q^2{}_2$, hence $x^1{}_1 = Hq^2{}_1 < q^2{}_1 < q^2{}_2 = x^1{}_2$, hence by (5) $p^1 = (0,1)$.

(c) Now define $\gamma^i = f^i((1,0),(1/2,1/2))_1$. By (3) $\gamma^i > 1/2$ for i=1,2.

Let $h = \max\{[\gamma^1]^{-1},[\gamma^2]^{-1}\}/2$, and note that $h < 1$. Consider $H > h$.

With χ defined as in section III, let $\chi^o = \{x^* \in \chi$: for i=1,2 $x^i{}_1 \geq 1/2$ and $x^i{}_2 \leq 1/2\}$ (note $\chi^o \neq \phi$).

Next define $T_{(1,0)}: M \times \chi^o \rightarrow M \times \chi$ as follows: if $T_{(1,0)}(\underline{q},x^*) = (\underline{s},y^*)$ then for i=1,2

$$s^i = f^i((1,0),x^i) \quad \text{and} \quad y^i = (Hs^{J(i)}{}_1, s^{J(i)}{}_2) \quad \text{where} \quad J(i) = \begin{cases} 2 & \text{if } i=1 \\ 1 & \text{if } i=2. \end{cases}$$

Then assumption (5) implies $s^i{}_1 \geq \gamma^i > 1/2$ and $s^i{}_2 < 1/2$, hence $y^i{}_1 = Hs^{J(i)}{}_1 \geq H\gamma^{J(i)} > h\gamma^{J(i)} \geq 1/2$ and $y^i{}_2 = s^{J(i)}{}_2 < 1/2$.

It follows that range $T_{(1,0)} \subseteq M \times \chi^o$ and hence $T_{(1,0)}$ has a fixed point, say $(\underline{q},x^*)$, where necessarily $(((1,0),(1,0)),\underline{q},x^*)$ is an imperfect equilibrium for $\Gamma(H)$.

In a similar fashion, define $h' = \min\{[f^1((0,1),(1/2,1/2))_1]^{-1},[f^2((0,1),(1/2,1/2))_1]^{-1}\}$, noting that $h' > 1$. Considering $H < h'$, we define $\chi' = \{x^* \in \chi$: for i=1,2 $x^i{}_1 \leq 1/2$ and $x^i{}_2 \geq 1/2\}$ and then $T_{(0,1)}: M \times \chi' \rightarrow M \times \chi'$ is defined analogously to $T_{(1,0)}$ above. $T_{(0,1)}$ must have a fixed point, say $(\underline{q},x^*)$, and $(((0,1),(0,1)),\underline{q},x^*)$ is an imperfect equilibrium for $\Gamma(H)$.

Without repeating the details, we conclude that for h<H<h' both ((1,0),(1,0)) and ((0,1),(0,1)) are target equilibria for $\Gamma(H)$.

APPENDIX B: PROOF OF PROPOSITION FOR PRISONER'S DILEMMA GAME

In $\Gamma(H)$ each of the two players has eight pure strategies, each of which is identified by a triple $(\sigma_1;\sigma_2,\sigma_3)$ where $\sigma_i \in \{c,d\}$, to be interpreted as: σ_1 is used in the first stage; subsequently, if the other party cooperates in the first stage then σ_2 is used at the second stage and if the other party defects in the first stage then σ_3 is used in the second stage. These strategies are renumbered 1-8 as follows.

1. c;cc	5. d;cc
2. c;cd	6. d;cd
3. c;dc	7. d;dc
4. c;dd	8. d;dd

Suppose that (p,q,x*) is an imperfect equilibrium for $\Gamma(H)$. For convenience, we shall use a short notation for the following expressions.

$$Q^i_c=q^i_1+q^i_2+q^i_3+q^i_4 \qquad Q^i_d=q^i_5+q^i_6+q^i_7+q^i_8$$

$$\Delta^i=(q^i_2+q^i_6)-(q^i_3+q^i_7)$$

These entities have a very intuitive meaning: Q^i_c and Q^i_d are, respectively, the probabilities of cooperating and of defecting in the first iteration, and Δ^i is the difference between the probabilities of using the strategies (2)-(6), which imitate what the opposition does, and those of using the strategies (3)-(7), which react in the reverse manner.

Also, instead of studying x* it will be convenient to look at the differences $x^i_8-x^i_k$ for k=1,...7. Straightforward computation gives the following, where (as before) J(1)=2 and J(2)=1.

$$x^i_8-x^i_1=1-\Delta^{J(i)}H+1 \qquad x^i_8-x^i_5=1$$

$$x^i_8-x^i_2=1-\Delta^{J(i)}H+Q^{J(i)}_c \qquad x^i_8-x^i_6=Q^{J(i)}_c$$

$$x^i_8-x^i_3=1-\Delta^{J(i)}H+Q^{J(i)}_d \qquad x^i_8-x^i_7=Q^{J(i)}_d$$

$$x^i_8-x^i_4=1-\Delta^{J(i)}H \qquad x^i_8-x^i_8=0$$

Note that either $\mu(x^i)=x^i_8$ or $\mu(x^i)=x^i_4$ (possibly both).

The analysis of the proposition begins with the observation that assumption (1) and (3) imply that the probability of an inadvertent use of a pure strategy involving an opportunity loss depends only on this opportunity loss and on the player's capability, i.e. for every player i there is a function $g^i:\Re_+\rightarrow[0,1]$ such that

if $p^i_k=0$ and $x^i_k<\mu(x^i)$ then $\qquad f^i(p^i,x^i)_k=g^i(\mu(x^i)-x^i_k)$.

Let δ be a given, arbitrarily small, positive number. Then by (6) $g^i(1/2-\delta)-g^i(1/2+\delta)>0$. Now define

$$h^*(\delta)=\max\{[g^1(1/2-\delta)-g^1(1/2+\delta)]^{-1},[g^2(1/2-\delta)-g^2(1/2+\delta)]^{-1}\}$$

Suppose $Q^2_c<1/2-\delta$, and thus $Q^2_d>1/2+\delta$. Then $x^1_7<x^1_6$, hence $p^1_7=0$ and by (3) and (6) $q^1_7<q^1_6$. Also $x^1_3<x^1_2$, similarly implying $q^1_3<q^1_2$. From the table of $x^i_x-x^i_8$ above we then have:

for either k=8 or k=4 $\qquad$ (i) $\mu(x^1)=x^1_k$

(ii) $f^1(p^1,x^1)_{k-1}=g^1(Q^2_d)<g^1(1/2+\delta)$

(iii) $f^1(p^1,x^1)_{k-2}=g^1(Q^2_c)>g^1(1/2+\delta)$.

Hence $\Delta^1>g^1(1/2-\delta)-g^1(1/2+\delta)$. Now if $H>h^*(\delta)$ then $\Delta^1H>1$,

hence for $k=1,\ldots,4$ $\; x^2_{k+4}-x^2_k=1-\Delta^1H<0$, implying $p^2_{k+4}=0$ with $q^2_{k+4}\leq q^2_k$.

Hence $Q^2_c\geq Q^2_d$ - a contradiction of the initial supposition $Q^2_c<1/2-\delta$.

Alternatively, suppose $Q^2_c>1/2$, with $Q^2_d<1/2$. Then $x^1_7>x^1_6$, with $p^1_6=0$ and $q^1_7>q^1_6$, and $x^1_3>x^1_2$ with $q^1_3>q^1_2$, together implying $\Delta^1<0$.

Hence for $k=1,\ldots,4$ $\; x^2_{k+4}>x^2_k$, $p^2_k=0$, and $q^2_k<q^2_{k+4}$, implying $Q^2_c<Q^2_d$ - again a contradiction.

It follows that $1/2-\delta\leq Q^2_c\leq 1/2$ whenever $H>h^*(\delta)$.

An identical argument applies for Q^1_c, completing the proof of the proposition.

REFERENCES

1. Aumann, Robert J. and Sylvain Solrin (1989), "Cooperation and Bounded Recall", *Games and Economic Behavior* **1**, 5-39.

2. Fudenberg, Drew and David K. Levine (1989), "Reputation and Equilibrium Selection in Games with a Patient Player", *Econometrica* **57**, 759-778.

3. Fudenberg, Drew and Eric Maskin (1986), "The Folk Theorem for Repeated Games with Discounting and with Incomplete Information", *Econometrica* **54**, 533-554.

4. Gilboa, Itzhak and Dov Samet (1989), "Bounded Versus Unbounded Rationality: The Tyranny of the Weak", *Games and Economic Behavior* **1**, 213-221.

5. Kreps, David, Paul Milgrom, John Roberts and Robert Wilson (1982), "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma", *Journal of Economic Theory* **27**, 245-252.

6. Kreps, David and Robert Wilson (1982), "Reputation and Imperfect Information", *Journal of Economic Theory* **27**, 253-279.

7. Milgrom, Paul and John Roberts (1982), "Predation, Reputation, and Entry Deterence", *Econometrica* **50**,443-460.

8. Myerson, Roger B. (1978), "Refinement of the Nash Equilibrium Concept", *International Journal of Game Theory* **7**, 73-80.

9. Neyman, Abraham (1985), "Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner's Dilemma", *Economic Letteres* **19**, 227-230.

10. Rubinstein, Ariel (1986), "Finite Automata Play the Repeated Prisoner's Dilemma", *Journal of Economic Theory* **39**, 83-96.

11. Selten, Reinhard (1975), "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games", *International Journal of Game Theory* **4**, 25-55.

12. Selten, Reinhard (1977), "The Chain-Store Paradox", *Theory and Decision* **9**, 127-259.

Date Due